

Thermal blinding of gated detectors in quantum cryptography

Lars Lydersen,^{1,2,*} Carlos Wiechers,^{3,4,5} Christoffer Wittmann,^{3,4}

Dominique Elser,^{3,4} Johannes Skaar,^{1,2} and Vadim Makarov¹

¹*Department of Electronics and Telecommunications,
Norwegian University of Science and Technology, NO-7491 Trondheim, Norway*

²*University Graduate Center, NO-2027 Kjeller, Norway*

³*Max Planck Institute for the Science of Light, Günther-Scharowsky-Str. 1/Bau 24, 91058 Erlangen, Germany*

⁴*Institut für Optik, Information und Photonik, University of*

Erlangen-Nuremberg, Staudtstraße 7/B2, 91058 Erlangen, Germany

⁵*Departamento de Física, Universidad de Guanajuato, Lomas del Bosque 103,
Fraccionamiento Lomas del Campestre, 37150, León, Guanajuato, México*

(Dated: September 15, 2010)

It has previously been shown that the gated detectors of two commercially available quantum key distribution (QKD) systems are blindable and controllable by an eavesdropper using continuous-wave illumination and short bright trigger pulses, manipulating voltages in the circuit [L. Lydersen *et al.*, Nat. Photonics DOI:10.1038/nphoton.2010.214]. This allows for an attack eavesdropping the full raw and secret key without increasing the quantum bit error rate (QBER). Here we show how thermal effects in detectors under bright illumination can lead to the same outcome. We demonstrate that the detectors in a commercial QKD system Clavis2 can be blinded by heating the avalanche photo diodes (APDs) using bright illumination, so-called *thermal blinding*. Further, the detectors can be triggered using short bright pulses once they are blind. For systems with pauses between packet transmission such as the plug-and-play systems, thermal inertia enables Eve to apply the bright blinding illumination *before* eavesdropping, making her more difficult to catch.

PACS numbers: 03.67.Dd

I. INTRODUCTION

In theory quantum mechanics allows two parties, Alice and Bob, to grow a private, secret key, even if the eavesdropper Eve can do anything permitted by the laws of nature [1–4]. The field of quantum key distribution (QKD) has evolved rapidly in the last two decades, with transmission distance increasing from a table top demonstration to over 250 km in the laboratory [5], and commercial QKD systems available from several vendors [6].

However the components used for the experimental realizations of QKD have imperfections. Numerous imperfections have been addressed in security proofs [7–12]. For some loopholes it took several years from their discovery until they were covered by security proofs, for instance the Trojan-horse [13, 14] loophole and detector efficiency mismatch [15, 16]. The latter was exploited in the time-shift attack [17] on a commercial QKD system [18]. Other loopholes include a variety of side-channels [19–22].

Common to the loopholes mentioned so far is that they are not implementable in practice, or only leave a marginal advantage for Eve. For instance, the implementation of the time-shift attack [18] gave Eve an information-theoretic advantage, allowing her to outperform a straight brute-force search for the key in 4% of her attempts. In the practical phase-remapping attack [22], Eve caused 19.7% QBER compromising merely the

hardly ever used two-way post-processing protocol which produces secure key at QBER up to 20% [23, 24].

There is however one class of attacks which stands out in terms of implementability, Eve’s information and QBER: The *blinding attacks* [25–27] are fully implementable with current technology, and give Eve the whole raw key while causing zero additional QBER. In these attacks, the APDs are tricked to exit the single-photon sensitive Geiger mode, and are so-called *blind*. Eve uses a copy of Bob’s apparatus to detect Alice’s signals, but resends bright trigger pulses instead of single photons, as in the after-gate attack [28]. When the detectors are blind, Bob will only detect the bright trigger pulses if he uses the same basis as Eve. Otherwise his detectors remain silent. Hence Eve gets a full copy of the raw key while causing no additional QBER. Both passively quenched detectors [25], actively quenched detectors [26] and the gated detectors of two commercially available QKD systems [27] have been shown to be vulnerable to blinding. In the case of the passively-quenched detectors, this loophole has been exploited in the first full-scale implementation of an eavesdropper [29], which was inserted in the middle of the 290 m transmission line in an experimental entanglement-based QKD system [30, 31], and recovered 100% of the raw key.

Previously the gated detectors in the commercially available system Clavis2 from manufacturer ID Quantique were subject to continuous-wave (CW) blinding [27]. The blinding illumination caused the bias voltage at the APDs to drop due to the presence of DC impedance of the bias voltage supply, and therefore the APDs were never in Geiger mode. In this paper we show how the

* lars.lydersen@iet.ntnu.no

same detectors, regardless of the impedance of the bias voltage supply, can be blinded by heating the APD, so-called *thermal blinding*. We show that thermal blinding is more sophisticated form of attack than previously reported CW-blinding [27] because the APD can be heated well in advance of the detection times, and is as such harder to catch. Especially for Clavis2, all the detector parameters such as temperature of the cold plate, bias voltage and APD current indicate single photon sensitivity while the detectors are in fact blind.

In this paper we first briefly review how APDs in the linear mode can be exploited to eavesdrop on QKD systems (section II). Then the detector design in Clavis2 is discussed (section III) before we show how it is possible to thermally blind and trigger the detectors (section IV). Finally we briefly discuss countermeasures in section V and conclude in section VI.

II. EAVESDROPPING EXPLOITING APDs IN LINEAR MODE

In this section we briefly review how APDs in the linear mode can be exploited to eavesdrop on QKD systems [26, 27].

In Geiger mode operation, an electron-hole pair produced by an absorbed single photon is amplified to a large current in the APD, which exceeds a current comparator threshold and reveals the photon's presence. This is referred to as a *click* [32].

In the linear mode however, when an APD is reverse-biased at a constant voltage below the breakdown voltage [33], the current through the APD is proportional to the incident optical power. Usually the APD is placed in a resistive network, and also has an internal resistance. Hence, the current through the APD lowers the bias voltage, and the current through the APD is monotonically increasing with the incident optical power. In this regime, the comparator current threshold translates to a classical optical power threshold [27].

If APDs are used as detectors in a QKD system, and they are optically accessible to Eve when biased under the breakdown voltage, Eve may eavesdrop on the QKD system with an intercept-resend (faked-state [34]) attack. Eve uses a copy of Bob to detect the qubits from Alice in a random basis. Eve resends her detection results, but instead of sending single photons she sends bright pulses, just above the classical optical power threshold. Bob will only have a detection event if his basis choice coincides with Eve's basis choice (see Fig. 1), otherwise no detector clicks.

After the raw key exchange, Bob and Eve are identical both in bit values and basis choices. Since Eve uses a copy of Bob's detectors, Bob's photon-number detection statistics is equal with or without Eve. Therefore the attack works equally well on the BB84 protocol [1], the Scarani-Acin-Ribordy-Gisin 2004 (SARG04) [35] and decoy-state BB84 protocols [36–38]. In addition to at-

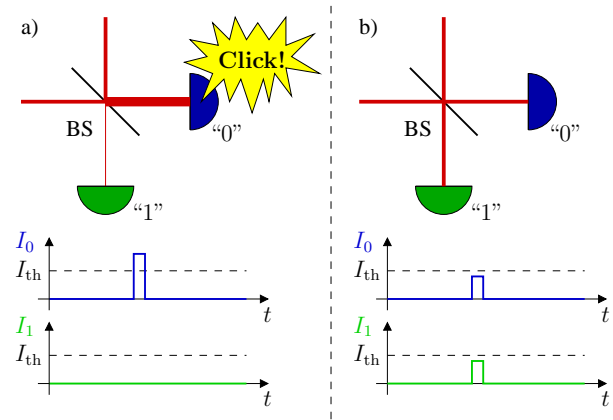


FIG. 1. The last beam splitter (BS) as well as the detectors in a phase-encoded QKD system. I_0 and I_1 is the current running through APD 0/1, and I_{th} is the comparator threshold current above which the detector registers a click. Here we assume that the APDs are in the linear mode, and that Eve sends a bright pulse slightly above the optical power thresholds. a) Eve and Bob have selected matching bases. Therefore the full intensity in the pulse from Eve hits detector 0. The current caused by Eve's pulse crosses the threshold current and causes a click. b) Eve and Bob have selected opposite bases. Therefore half the intensity of Eve's pulse hits each detector (corresponding to 50% detection probability in either detector for single photons). This causes no click as the current is below the threshold for each detector.

tacking the quantum channel, Eve listens on the classical channel between Alice and Bob. Afterwards Eve performs the same classical post-processing as Bob to obtain the identical secret key.

Note that the classical optical power threshold has to be sufficiently well defined for successful perfect eavesdropping. To be precise, let an optical power of $P_{100\%,i}$ or greater always cause a click when applied to detector i . Likewise, let an optical power of $P_{0\%,i}$ or less never cause a click when applied to detector i . The sufficient condition for Eve to be able to make any single detector click while none of the other detectors click, can be expressed as

$$\max_i \{P_{100\%,i}\} < 2 \left(\min_i \{P_{0\%,i}\} \right). \quad (1)$$

III. DETECTOR DESIGN

A. Detector circuit

Figure 2 shows an equivalent detector bias and comparator circuit diagram for the detectors in Clavis2, based on reverse engineering. The APD is biased just above its breakdown voltage by the high voltage supply $V_{HV,0} = -42.89$ V, $V_{HV,1} = -43.08$ V. On top of this bias the APD is gated with 2.8 ns TTL pulses every 200 ns from DD1 to create Geiger mode gates. The gates

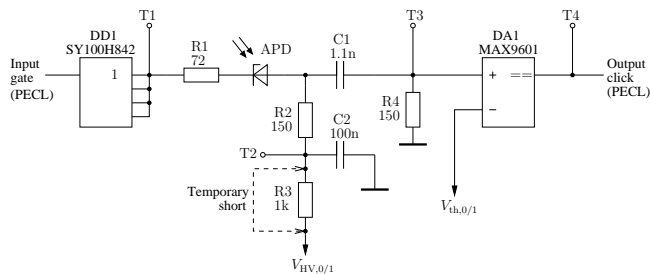


FIG. 2. Equivalent detector bias and comparator circuit. Taps T1-T3 are analog taps of the APD gates ($V_{\text{gate},0/1}$), the APD bias ($V_{\text{bias},0/1}$) and the comparator input ($V_{\text{comp},0/1}$). The digital tap T4 of the detector output ($V_{\text{click},0/1}$) has been converted to logic levels in all oscillograms. For the experiments presented in section IV, the resistor R3 has been shorted.

are applied as PECL signals from the mainboard, and the buffer converts them to TTL levels, 0 V and approximately 3 V. The anode of the APD is AC-coupled to a fast comparator DA1 with the thresholds $V_{\text{th},0} = 78$ mV and $V_{\text{th},1} = 82$ mV.

The normal operation of the detector circuit can be seen in Fig. 3. A number of techniques have been developed for compensating the capacitive pulse through APDs in the absence of an avalanche [39–42], but this particular detector simply sets the comparator thresholds above the amplitude of the capacitive pulse.

As a side note, applying CW illumination to the APD allowed us to measure the timing of the quantum efficiency curve within the gate quite precisely, see Appendix B.

B. Detector cooling

To reduce the probability of dark counts, APDs are usually cooled to a low temperature. The two APDs in this QKD system are cooled together by one 4-stage thermoelectric cooler (TEC) (Ostern PE4-115-14-15 [43]). The system software reports the temperature measured by a thermistor mounted on the cold side of the top stage (cold plate), and close to where the APDs are mounted. Note that the cold plate temperature is not always the same as the APD chip temperature, as there is actually a quite substantial thermal resistance between the two. This will become an important point in section IV B. The hot side of the TEC is mounted on a large heatsink with a fan, such that it stays at approximately room temperature.

The temperature of the cold plate is maintained at a pre-set value by a closed-loop controller that adjusts the TEC current. When the system is switched on, the cold plate (and thus the APDs) is first cooled to the target temperature, -50°C . The system will not start operation unless the cold plate settles at a temperature below -49.8°C . After this initial check however, during system




FIG. 3. An example of electrical signals during two gates in detector 1 without any illumination. In the first gate thermal fluctuations or trapped carriers have caused an avalanche, and a click at the comparator output (dark count). A typical amplitude of the avalanche peak is 200 mV for detector 0 and 300 mV for detector 1. Normally the system removes 50 gates after a detection event, but for this oscillogram this feature has been disabled. In the second gate there is no detection event. When no current runs through the APD, it is equivalent to a capacitor, and thus approximately the derivative of the gate pulse shape propagates to the comparator input, with peak positive amplitude ≈ 35 mV.

operation, there seems to be no future checks of the cold plate temperature, even if the controller is unable to keep it at the target value.

IV. BLINDING AND CONTROL

Blinding is achieved when the system is insensitive to single photons. This can be achieved by ensuring that the APD bias voltage is below the breakdown voltage, or by lowering the voltage in front of the comparator such that the avalanche current does not cross the comparator threshold. The detectors are controllable if they are accessible to Eve in the linear mode with a sufficiently well defined classical optical power click threshold, as in Eq. 1.

We have previously reported that blinding Clavis2 can be achieved by CW illumination due to the bias voltage supply impedance $R3 = 1$ k Ω , which makes the bias voltage drop to a level where the APD is never in Geiger mode [27], even inside the gate.

One fast and easy countermeasure could be to use a low-impedance bias voltage source in the detectors. Therefore, in this paper we consider a modified version of the detectors with R3 shorted (see Fig. 2). We present three different blinding techniques which may be used against detectors with a low-impedance bias voltage source, and show that the detectors can be controlled by trigger pulses in the blind state. The technique in section IV A clearly works against high-impedance biased

detectors as well as against low-impedance biased detectors since it has been demonstrated [27]. The difference is that with a low-impedance bias voltage source, the blinding originates from thermal effects instead of bias voltage drop. The technique in section IV B has been used on low-impedance biased detectors, but we see no reason why it should not work similarly well against the unmodified high-impedance biased detectors. The technique in section IV C has been used on both high- and low-impedance biased detectors, but we only present the results for the low-impedance biased detectors in this paper.

A. Thermal CW-blinding

It turns out that it is possible to blind also low-impedance biased detectors ($R3 = 0$) by CW illumination. When an APD is illuminated, the power dissipated in the APD is transformed to heat, which may increase the APD temperature. The breakdown voltage is temperature dependent: increasing the temperature increases the breakdown voltage. Since the bias voltage is constant, this makes the APD leave the Geiger mode. Two effects contribute to the power dissipation: electrical heating ($V_{APD} \cdot I_{APD}$) and the small contribution by the absorption of the optical power. For the heat dissipation calculations, we simply assume that all the optical power is absorbed and transformed to heat. Figure 4 shows how the heat dissipation increases with the optical illumination.

When the sum of the heat dissipations of the two detectors is approximately 300 mW, the cooling system is running at its maximum capacity with a TEC current of about $I_{TEC} = 2.37$ A (the air temperature at the heatsink fan intake at this time was 23.6 °C). When the optical illumination is increased beyond this point, the cold plate (and thus APD) temperature starts to increase. Figure 5 shows how the temperature of the cold plate increases with the total amount of heat dissipated in the APDs. When the optical illumination, and thus the load is increased beyond the maximum capacity of the TEC, the cold plate temperature increases approximately linearly with the heat dissipated by the APD. While not in the specifications of this specific TEC [43], other data sheets of similar TECs [44] show that the temperature difference between the hot and cold plate decreases linearly with respect to the load, given a constant TEC current.

When the temperature of the APDs increases, the breakdown voltage also increases with the coefficient of about 0.1 V/K [45]. In this experiment we illuminated both detectors simultaneously, to get sufficient temperature increase without risking a permanent damage to the APDs. We used a fibre-optic coupler (see appendix A for the experimental setup) to illuminate both detectors, with 46.75%/53.25% of the optical power going to detector 0/1. This is approximately equal to the measured splitting ratio for the beam splitter in front of the detec-




FIG. 4. Calculated heat dissipation (based on measured APD current and voltage) versus the optical illumination for each of the two detectors.




FIG. 5. The temperature of the cold plate and TEC current reported by the software, versus the total amount of heat dissipated in the APDs. It takes several minutes for the cold plate temperature to stabilize at a new value (hotter than -50 °C) after the power dissipation in the APDs is changed.

tors in the system, when illuminated through the short arm of the interferometer [46–48].

Fig. 6 shows the click probability versus the CW illumination of the two detectors. The click probability drops below the normal dark count probability (about 10^{-4}), before it becomes *exactly zero* when the illumination exceeds 8.8 mW and 10 mW at the detectors. In the experiment the blinding caused clicks for several minutes before the APDs were properly heated. However, the blinding only needs to be turned on once, afterwards Eve remains undetected.

After the cold plate has been heated by APD illumination, it takes several tens of seconds before it cools to the target temperature of -50 °C. Therefore, the detectors stay blind for some time after the CW blinding illumination is turned off. Detectors 0 and 1 regain dark counts when the cold plate (and thus the APDs) becomes colder than -39.8 °C and -40.1 °C, respectively.

To verify that the detectors could be controlled, the detectors were blinded with 9.5 mW at detector 0 and 10.7 mW at detector 1, and controlled by superimposing




FIG. 6. Click probability versus power of CW illumination applied to both detectors simultaneously.

TABLE I. Control pulse peak power at 0 % and 100 % click probability thresholds, in CW thermal blinding mode.

Detector	Click probabilities	
	0 %	100 %
0	1.12 mW	1.31 mW
1	1.71 mW	2.02 mW

a 3 ns long laser pulse slightly after the gate. The click probability thresholds are listed in table I. The thresholds satisfy Eq. 1, and thus the eavesdropping method described in section II should be possible when the detectors are thermally blinded by CW illumination.

After observing thermal blinding in this experiment, we realized that this could be the reason why the PerkinElmer SPCM-AQR actively-quenched detector module remained blind at bright pulse frequencies above 400 kHz, despite no substantial bias voltage drop [26]. Therefore we did more precise measurements which confirm that PerkinElmer SPCM-AQR can be thermally blinded [49].

B. Thermal blinding of frames

As this QKD system is of plug-and-play type, it sends the qubits in packets called *frames* to avoid Rayleigh back-scattered photons to arrive during the gates and increase the QBER [46, 50]. For our experiment we used 1072 qubits per frame [51]. With a 200 ns bit period this makes the frame length 214.4 μ s. The break in between the frames varies with the fibre length between Alice and Bob, but is always longer than the frame itself. In our experiment we simply used a 250 μ s frame break, which makes a total frame + break period of 464.4 μ s.

It turns out that the APD chip and the inner parts immediately touching it (*not* the APD package and not the cold plate) act as a thermal reservoir on the frame period time scale. Therefore bright illumination between the frames heats the APD sufficiently that it stays blind throughout the whole frame. Based on the optical power

where the frames went blind, and the average current through the APDs, the thermal resistance between each APD chip and the cold plate is estimated to be at least 190 K/W.

To heat the APDs we used 225 μ s long pulses timed in between the frames and fired at both APDs simultaneously. The whole frame went blind at approximately 1.5 mW and 1.7 mW pulse power at detector 0 and 1 respectively. The oscillograms in Fig. 7 show the electrical and optical signals in detector 1 when frames of 1072 gates are thermally blinded by the 225 μ s long pulses with 3.5 mW in-pulse power at detector 0, and 4 mW in-pulse power at detector 1. While the system was blind, the cold plate temperature reading was -49.5°C , and the TEC was running well below its maximum capacity at $I_{\text{TEC}} = 2.006$ A.

To verify that the detectors could be controlled, we checked the response to a 4 ns long control pulse timed slightly after the gate of one of the first bits of the frame, and the last bit of the frame. The detection probability thresholds for the second [52] and the last bit are given in tables II and III. Figure 8 shows oscillograms from detector 1 when it is blinded and controlled in the second bit of the frame.

The click probability thresholds in tables II and III each satisfy Eq. 1 individually. However, $P_{0\%,0}$ in the last bit of the frame is less than $1/2$ of $P_{100\%,1}$ in the second bit of the frame. This means that the control pulse power would have to be decreased throughout the frame. Since the second and the last bit of the frame can be controlled, it is plausible that the eavesdropping method described in section II could be applied to any bit of the frame.

What is remarkable about this blinding method is that due to the low thermal conductivity between the APD chip and the cold plate, as well as the thermal inertia of the nearby parts, the cold plate thermistor reports a value very close to the normal value. Therefore monitoring the cold plate temperature would not suffice to prevent thermal blinding.

TABLE II. Control pulse peak power at 0 % and 100 % click probability thresholds for the second bit in the frame, when the frame is thermally blinded.

Detector	Click probabilities	
	0 %	100 %
0	401 μ W	533 μ W
1	580 μ W	747 μ W

TABLE III. Control pulse peak power at 0 % and 100 % click probability thresholds for the last bit in the frame, when the frame is thermally blinded.

Detector	Click probabilities	
	0 %	100 %
0	305 μ W	420 μ W
1	340 μ W	532 μ W




FIG. 7. Thermal blinding of frames. The oscillograms show electrical and optical signals when frames of 1072 gates in detector 1 are thermally blinded by a 225 μs blinding pulse, with 3.5 mW pulse power at detector 0, and 4 mW pulse power at detector 1. The blinding pulse causes a detection event outside the frame, where the system probably does not register clicks (If the click is registered, it could easily be avoided by increasing the power of the blinding pulse gradually, such that the comparator input AC-coupling keeps the voltage below the comparator threshold).

C. Sinkhole blinding

It is natural to ask whether the framed blinding technique can be applied at the single gate level, i.e. what happens if bright illumination is applied between adjacent gates? It turns out that this also leads to blinding, but not primarily due to thermal effects. Since the comparator input is AC-coupled (see Fig. 2), the signal at the input of the comparator has the same area over and un-

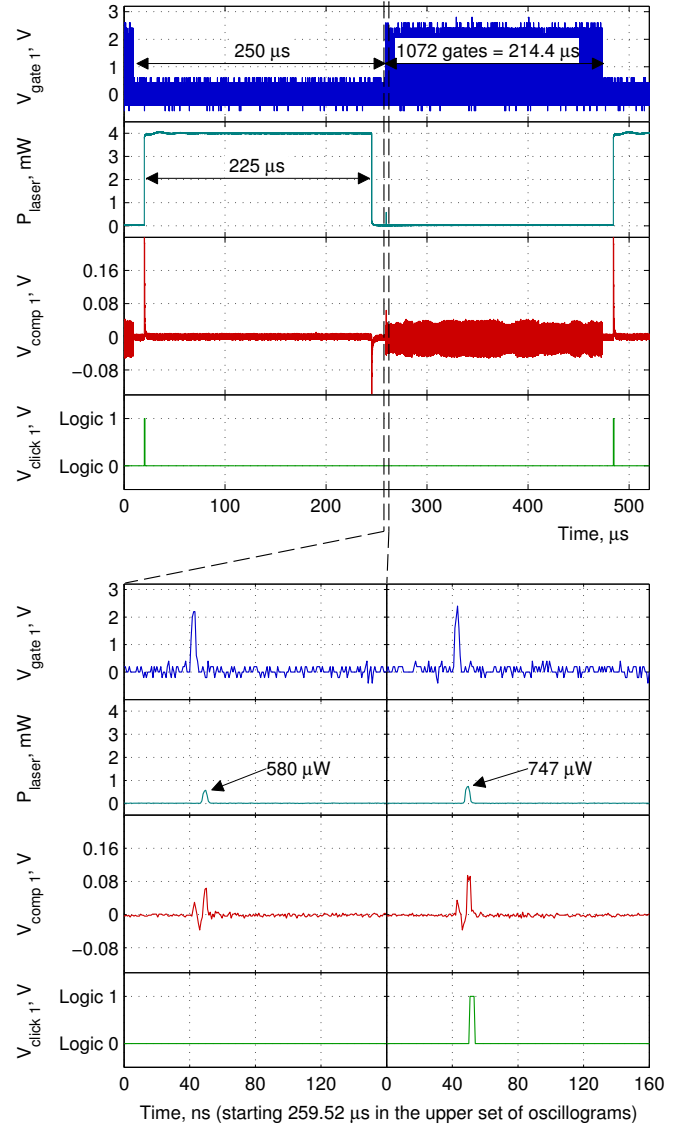


FIG. 8. Detector control during thermal blinding of frames. The oscillograms show electrical and optical signals when frames of 1072 gates in detector 1 are thermally blinded by a 225 μs blinding pulse, with 3.5 mW pulse power at detector 0, and 4 mW pulse power at detector 1, and the detector is controlled by a 4 ns long control pulse timed slightly after the second gate in the frame. In the upper and lower left sets of oscillograms, the 580 μW control pulse never causes any click. In the lower right set, the control pulse is applied after the same gate in the frame, but now its increased 747 μW peak power always causes a click.

der 0 V level when averaged over time much longer than $R_4 \cdot C_1 = 165 \text{ ns}$. Thus by sending long bright pulses between the gates and no illumination near the gate, it is possible to superimpose a negative-voltage pulse at the comparator input at the gate time. We call this negative pulse a *sinkhole*. An avalanche that occurs within it can have a normal amplitude yet remain below the comparator threshold level.




FIG. 9. Sinkhole blinding. The oscillograms show electrical and optical signals when detector 1 is blinded by a $500 \mu\text{W}$, 140 ns long laser pulse in between the gates. The avalanche amplitude is about 130 mV and would cause a click if it were not sitting in the negative-voltage pulse. It seems that the reduction in avalanche amplitude (compare to Fig. 3) is caused by heating of the APD, which effectively rises the breakdown voltage.

TABLE IV. Control pulse peak power at 0 % and 100 % click probability thresholds, during sinkhole blinding.

Detector	Click probabilities	
	0 %	100 %
0	$655 \mu\text{W}$	$751 \mu\text{W}$
1	$773 \mu\text{W}$	$908 \mu\text{W}$

Using a 140 ns long pulse beginning about 25 ns after the gate, detector 0 becomes completely blind when $P_{\text{laser}} > 205 \mu\text{W}$, and detector 1 becomes blind when $P_{\text{laser}} > 400 \mu\text{W}$. To keep both detectors blind, $P_{\text{laser}} = 500 \mu\text{W}$ is used subsequently. When a large pulse is applied between the gates, the detector will always experience a dark count in the gate due to trapped carriers. Figure 9 shows detector 1 blinded by a 140 ns long, $500 \mu\text{W}$ bright pulse, starting about 25 ns after the gate.

Initially when the blinding pulses are turned on, there is a transient with about 20-100 clicks, which would be easily detectable in post-processing. Note again that the blinding only needs to be turned on once, and that the blinding can be turned on before the raw key exchange to avoid the clicks being registered.

Detector control is obtained by a 3.2 ns long laser pulse timed shortly after the gate. The click probability thresholds found are listed in Table IV. Figure 10 shows oscillograms from detector 1 when it is blind and controlled. Once again, the thresholds in table IV satisfy




FIG. 10. Detector control during sinkhole blinding. The oscillograms show electrical and optical signals when detector 1 is blinded with a $500 \mu\text{W}$, 140 ns long laser pulse in between the gates, and controlled with a 3.2 ns long laser pulse timed shortly after the gate. To the left, the $773 \mu\text{W}$ control pulse never causes any click. To the right, the $908 \mu\text{W}$ control pulse always causes a click.

Eq. 1, and thus the eavesdropping method described in section II should be possible when the detectors are sinkhole blinded.

V. DISCUSSION AND COUNTERMEASURES

First of all, the numerous detectors proved blindable and controllable [25–27, 29, 49], and the large number of independent blinding methods available show that avoiding this loophole is non-trivial. Further the results presented in this paper clearly show that reducing the impedance of the bias voltage supply is far from being a sufficient countermeasure for this detector design.

At this point it is not clear to us how to design hack-proof detectors. The most obvious countermeasure is to monitor the optical power at Bob's entrance with an additional detector. However it is not obvious that this actually closes the loophole; as pointed out previously the click threshold close to the gate may be very low, allowing for practically non-detectable control pulses [27]. Thus it is not clear how to set the threshold value for the entrance monitor; in any case the threshold should be derived from and incorporated into a security proof. It would also be crucial that this monitoring detector is not blindable.

For the passively quenched scheme it has been proposed previously to monitor APD parameters such as APD bias voltage, current and temperature [25]. How-

ever, the results in section IV B show that normal APD parameters do not necessarily guarantee single photon sensitivity: for thermal blinding of frames all the APD parameters report normal values during the frames while the detectors are in fact blind.

It is worth emphasizing that the loophole opens when Eve drives the detectors into an abnormal operating regime, namely the linear mode. However, there are also quantum detectors which are actually designed to operate in linear mode. For example, homodyne detectors used in continuous-variable QKD [53, 54] are probably not susceptible to the described attack.

VI. CONCLUSION

The detectors in the Clavis2 QKD system have proved to be blindable by a variety of methods, even with a low-impedance bias voltage supply. Further, the detectors can always be controlled in the blind state. This allows eavesdropping on the QKD system, using the method described in section II. Since Eve may use an exact copy of Bob's system, no parameters currently available to Bob reveal Eve's presence. In practice, this should allow for perfect eavesdropping where Eve has an exact copy of Bob's raw key, and thus can extract the full secret key. The eavesdropping strategy described in section II has been implemented and used to capture 99.8% of the raw key in a 290 m experimental entanglement-based QKD system [29]. We see no practical difficulties implementing the same eavesdropper for this commercial QKD system, using off-the-shelf components. Actually we have proposed a plug-and-play eavesdropper scheme [27] for easy deployment.

Many detectors have already been proved blindable and controllable by Eve [25–27], and the large variety of blinding methods available for the system tested could probably be used on other detector designs as well. While it is relatively easy to design a countermeasure that prevents blinding attacks with the specific parameters chosen in the present work, it is unclear to us how to build generic secure detectors.

This work further emphasizes the importance of thoroughly investigating the non-idealities of each component in a QKD system, as well as battle-testing the system as a whole.

ID Quantique has been notified about the loophole prior to this publication, and has implemented counter-measures.

Appendix A: Measurement setup

Figure 11 shows the measurement setup used for this experiment. The trigger signal is tapped directly from the PECL gate signal (before DD1 in Fig. 2).

When pump current is used to control the power of the laser, the pulse width will vary slightly with the peak




FIG. 11. The setup used in the experiment. Both detectors were illuminated simultaneously by inserting a 50/50 fibre-optic coupler (not shown in the diagram) before the APDs.

power. In our experiment, the observed change in pulse width is less than 10 % after doubling the laser power. Also, the comparator threshold does not seem to be significantly dependent on the pulse width, thus we consider our results valid despite this small change in the laser pulse width.

Appendix B: Direct measurement of quantum efficiency

When CW illumination is applied to the APD, the applied electrical gate “propagates” to the comparator input. This might be caused by a change in linear multiplication coefficient caused by the electrical gate. This allowed us to measure the quantum efficiency mapped inside the “propagated” gate with about 200 ps precision.

The single photon sensitivity was measured using a id300 short-pulsed laser attenuated to a mean photon number of 1 per pulse. The quantum efficiency η was derived from the data assuming that the detector is linear (i.e. that an n -photon state is detected with probability $1 - (1 - \eta)^n$). The timing of the photon arrival at the APD relative to the applied gate was aligned by observing a response to unattenuated laser pulse on top of the 2.1 mW CW illumination. Figure 12 shows the result of the measurement on detector 1.




FIG. 12. Quantum efficiency measured directly within the electrical gate for detector 1. The photon sensitivity drops about 1 ns before the falling edge of the gate, because avalanches that start late do not have time to develop a large enough current to cross the comparator threshold.

ACKNOWLEDGMENTS

This work was supported by the Research Council of Norway (grant no. 180439/V30) and DAADppp mobil-

ity program financed by NFR (project no. 199854) and DAAD (project no. 50727598).

-
- [1] C. H. Bennett and G. Brassard, in *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing* (IEEE Press, New York, Bangalore, India, 1984) pp. 175–179.
 - [2] A. K. Ekert, *Phys. Rev. Lett.*, **67**, 661 (1991).
 - [3] H.-K. Lo and H. F. Chau, *Science*, **283**, 2050 (1999).
 - [4] P. W. Shor and J. Preskill, *Phys. Rev. Lett.*, **85**, 441 (2000).
 - [5] D. Stucki, N. Walenta, F. Vannel, R. T. Thew, N. Gisin, H. Zbinden, S. Gray, C. R. Towery, and S. Ten, *New J. Phys.*, **11**, 075003 (2009).
 - [6] Commercial QKD systems are available from at least two companies: ID Quantique (Switzerland), <http://www.idquantique.com>; MagiQ Technologies (USA), <http://www.magiqtech.com>.
 - [7] D. Mayers, in *Proceedings of Crypto'96*, Vol. 1109, edited by N. Kobitz (Springer, New York, 1996) pp. 343–357.
 - [8] D. Gottesman, H.-K. Lo, N. Lütkenhaus, and J. Preskill, *Quant. Inf. Comp.*, **4**, 325 (2004).
 - [9] H. Inamori, N. Lütkenhaus, and D. Mayers, *Eur. Phys. J. D*, **41**, 599 (2007).
 - [10] C.-H. F. Fung, K. Tamaki, B. Qi, H.-K. Lo, and X. Ma, *Quant. Inf. Comp.*, **9**, 131 (2009).
 - [11] L. Lydersen and J. Skaar, *Quant. Inf. Comp.*, **10**, 0060 (2010).
 - [12] Ø. Marøy, L. Lydersen, and J. Skaar, e-print arXiv:0903.3525, accepted for publication in *Phys. Rev. A*.
 - [13] A. Vakhitov, V. Makarov, and D. R. Hjølme, *J. Mod. Opt.*, **48**, 2023 (2001).
 - [14] N. Gisin, S. Fasel, B. Kraus, H. Zbinden, and G. Ribordy, *Phys. Rev. A*, **73**, 022320 (2006).
 - [15] V. Makarov, A. Anisimov, and J. Skaar, *Phys. Rev. A*, **74**, 022313 (2006), erratum *ibid.* **78**, 019905 (2008).
 - [16] V. Makarov and J. Skaar, *Quant. Inf. Comp.*, **8**, 0622 (2008).
 - [17] B. Qi, C.-H. F. Fung, H.-K. Lo, and X. Ma, *Quant. Inf. Comp.*, **7**, 73 (2007).
 - [18] Y. Zhao, C.-H. F. Fung, B. Qi, C. Chen, and H.-K. Lo, *Phys. Rev. A*, **78**, 042333 (2008).
 - [19] A. Lamas-Linares and C. Kurtsiefer, *Opt. Express*, **15**, 9388 (2007).
 - [20] S. Nauerth, M. Fürst, T. Schmitt-Manderbach, H. Weier, and H. Weinfurter, *New J. Phys.*, **11**, 065001 (2009).
 - [21] C.-H. F. Fung, B. Qi, K. Tamaki, and H.-K. Lo, *Phys. Rev. A*, **75**, 032314 (2007).
 - [22] F. Xu, B. Qi, and H.-K. Lo, e-print arXiv:1005.2376v1 [quant-ph].
 - [23] H. F. Chau, *Phys. Rev. A*, **66**, 060302 (2002).
 - [24] D. Gottesman and H.-K. Lo, *IEEE Trans. Inf. Theory*, **49**, 457 (2003).
 - [25] V. Makarov, *New J. Phys.*, **11**, 065003 (2009).
 - [26] V. Makarov, A. Anisimov, and S. Sauge, e-print arXiv:0809.3408v2 [quant-ph].
 - [27] L. Lydersen, C. Wiechers, C. Wittmann, D. Elser, J. Skaar, and V. Makarov, *Nat. Photonics* (2010), doi: 10.1038/nphoton.2010.214.
 - [28] C. Wiechers, L. Lydersen, C. Wittmann, D. Elser, J. Skaar, C. Marquardt, V. Makarov, and G. Leuchs, “After-gate attack on a quantum cryptosystem,” In preparation.
 - [29] I. Gerhardt, Q. Liu, J. Skaar, A. Lamas-Linares, C. Kurtsiefer, and V. Makarov, in preparation.
 - [30] I. Marcikic, A. Lamas-Linares, and C. Kurtsiefer, *Appl. Phys. Lett.*, **89**, 101122 (2006).
 - [31] M. P. Peloso, I. Gerhardt, C. Ho, A. Lamas-Linares, and C. Kurtsiefer, *New J. Phys.*, **11**, 045007 (2009).
 - [32] S. Cova, M. Ghioni, A. Lotito, I. Rech, and F. Zappa, *J. Mod. Opt.*, **51**, 1267 (2004).
 - [33] All references to the APD bias voltage are absolute valued, thus an APD biased “above” the breakdown voltage is in the Geiger mode. In practice the APDs are always reverse-biased.
 - [34] V. Makarov and D. R. Hjølme, *J. Mod. Opt.*, **52**, 691 (2005).
 - [35] V. Scarani, A. Acín, G. Ribordy, and N. Gisin, *Phys. Rev. Lett.*, **92**, 057901 (2004).
 - [36] W.-Y. Hwang, *Phys. Rev. Lett.*, **91**, 057901 (2003).
 - [37] X.-B. Wang, *Phys. Rev. A*, **72**, 012322 (2005).
 - [38] H.-K. Lo, X. Ma, and K. Chen, *Phys. Rev. Lett.*, **94**, 230504 (2005).
 - [39] S. Cova, A. Longoni, and A. Andreoni, *Rev. Sci. Instrum.*, **52**, 408 (1981).
 - [40] D. S. Bethune and W. P. Risk, *IEEE J. Quantum Electron.*, **36**, 340 (2000).
 - [41] A. Tomita and K. Nakamura, *Opt. Lett.*, **27**, 1827 (2002).
 - [42] Z. L. Yuan, B. E. Kardynal, A. W. Sharpe, and A. J. Shields, *Appl. Phys. Lett.*, **91**, 041114 (2007).
 - [43] Osterm, PE4-115-14-15, <http://osterm.ru/PAGE/MULTISTAGE.HTM>, visited 3. August 2010.
 - [44] Marlow, NL4012, <http://www.marlow.com/media/marlow/product/downloads/nl4012t/NL4012.pdf>, visited 3. August 2010.
 - [45] The detectors do not have any dark counts and are assumed blind at a temperature of about -40°C at the cold plate, or when the bias voltage is decreased by 0.97 V. If one assumes that the APD temperature is equal to the cold plate temperature, this means that heating the detectors by 10 K is equivalent to decreasing the bias voltage by about 1 V.
 - [46] G. Ribordy, J.-D. Gautier, N. Gisin, O. Guinnard, and H. Zbinden, *Electron. Lett.*, **34**, 2116 (1998).
 - [47] D. Stucki, N. Gisin, O. Guinnard, G. Ribordy, and H. Zbinden, *New J. Phys.*, **4**, 41 (2002).
 - [48] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, *Rev. Mod. Phys.*, **74**, 145 (2002).
 - [49] S. Sauge, L. Lydersen, A. Anisimov, J. Skaar, and V. Makarov, in preparation.
 - [50] G. Ribordy, J.-D. Gautier, N. Gisin, O. Guinnard, and H. Zbinden, *J. Mod. Opt.*, **47**, 517 (2000).

- [51] The system actually sends the qubits in frames of 1075 qubits each. We initially made a mistake when counting them and used 1072 qubits, which is very close and does not affect the results.
- [52] We picked the second bit to simplify synchronization in our measurement setup. The results for the first bit should be very similar to the results for the second bit.
- [53] S. L. Braunstein and P. van Loock, *Rev. Mod. Phys.*, **77**, 513 (2005).
- [54] U. L. Andersen, G. Leuchs, and C. Silberhorn, *Laser Photon. Rev.*, **4**, 337 (2010), arXiv:1008.3468v1 [quant-ph].